

### REMARKS

This request for reconsideration is responsive to the final Office Action dated August 22, 2006 and received in this application. Claims 1-11 and 18-26 remain pending in the application. Reconsideration of the pending claims in light of the following remarks is respectfully requested.

Claims 1-6, 9-11, 18-21, 23 and 25 have been rejected under 35 U.S.C. § 102 as being anticipated by U.S. Pat. No. 6,823,336 to Srinivasan et al. ("Srinivasan"). This rejection is traversed.

Claim 1 recites: *[a] system for mirroring write operations from a local storage system onto a remote storage system, the system comprising:*

*an asynchronous mirroring driver resident in the local storage system for intercepting I/O transactions to a storage disk of the local storage system, identifying a series of write transactions issued to said storage disk, making an exact copy of the series of write transactions, and storing said exact copy within a series of files that are created on a file-system of the local storage system; and*

*a first asynchronous mirroring coordinator resident on the local storage system for invoking a file transfer system to transmit the series of files on the local file-system of the local storage system to a file system of the remote storage system via a non-proprietary network protocol to accommodate an exact reproduction at the remote storage system of the series of write transactions as issued to said storage disk of the local storage system.*

Applicant's claimed invention provides asynchronous mirrored storage by intercepting I/O transactions to a storage disk of a local storage system, and retaining an exact copy of the corresponding write transactions within a series of regular file system files. The asynchronous mirroring coordinator may then transmit these files on any desired schedule to the file system of the remote storage system using a non-proprietary network protocol to accommodate an exact reproduction of the write transactions as issued to the storage disk of the local storage system.

As previously explained, these claimed features provide several advantages and distinctions over conventional systems. Retaining an exact record of write transactions allows a

return to any point on a per-transaction basis in the event of failure on either the local or remote storage side. The overhead of managing a local buffer and corresponding with the remote system in response to regular write transactions is also avoided. Finally, implementation of file system files and non-proprietary network communication protocols (*e.g.*, IP and/or FTP) introduces flexibility and resiliency to the system.

These claimed features are neither disclosed nor suggested by Srinivasan. Srinivasan discloses techniques for ensuring dataset consistency. This dataset consistency appears to pertain to high level datasets such as databases and files. Regardless, the Srinivasan system does not retain exact copies of write transactions, but instead implements a buffer, a directory of revisions, and corresponding read and write algorithms to service reads and writes while maintaining the “dataset” on the secondary system. For example, the Abstract states that:

“A data storage system receives sets of the revisions such that each set of revisions changes the dataset from one consistent state to another. Each set of revisions is processed in a write-selected phase followed by a read-selected phase. In the write-selected phase, the revisions in each set are written to a buffer and processed to produce a directory of the set of revisions. In the read-selected phase, the revisions are read from the buffer and integrated into the dataset. When one set of revisions is in the read-selected phase, the next set of revisions is in the write-selected phase. To permit uninterrupted read-only access to a consistent state of the dataset, the data storage system responds to a request for reading specified data on a priority basis by first accessing the directory of the read-selected revisions to determine whether the specified data are in the read-selected set of revisions, and if so, the specified data are read from the read-selected set of revisions, and if not, the specified data are read from the dataset.”

(Srinivasan, Abstract).

Srinivasan thus accumulates a “set of revisions” and then in some circumstances writes the revisions to a buffer, and in other circumstances reads the revisions from a buffer and integrates them into the dataset, while concurrently maintaining the dataset at the secondary location. Clearly, this scheme does not even generally set out to intercept and retain an exact copy of write transactions issued to a storage disk, or to store this exact copy in a file system file as claimed by Applicant. In that regard, at least the following claimed features are absent from Srinivasan: (1) the I/O transactions to a storage disk are intercepted and an exact copy of the write transactions is retained; (2) the retained exact copy of the write transactions is stored within a series of files created on the file system of the local storage system; and (3) the series of files is transmitted to the remote storage system using a non-proprietary network protocol to

accommodate an exact reproduction of the series of write transactions at the remote storage system.

First, Srinivasan relates to what is referred to as database consistency and thus seems to confront the problem of redundancy at a higher level rather than the storage level. Srinivasan maintains a “dataset” on the local and remote ends and employs buffers and different treatment of reads and writes to ensure that the dataset remains consistent on both the primary and secondary sides. This is not keeping an exact copy of the write transactions, and retaining them in a series of files created on the file system as claimed. At best, the dataset and corresponding techniques of Srinivasan could be abstractly construed as a scheme for consolidating writes, which is distinct from intercepting and retaining the exact copy of write transactions.

Additionally, Srinivasan’s updating of the dataset at the primary and secondary locations does not disclose in any way storing the exact copy of the write transactions in a series of file system files, and then transmitting the files containing the exact copy of the write transactions, as claimed by Applicant. The dataset of Srinivasan is, of course, stored in some fashion at the primary and secondary locations. However, this is simply a database, file, etc. that is kept consistent on both the local and remote sides using the scheme of keeping a set of revisions, using buffer(s) and managing reads and writes. There is no mention whatsoever of encapsulation of an exact copy of the write transactions within file system files as claimed.

Finally, Srinivasan merely mentions “transmission line 22” and does not specify a non-proprietary network protocol for transmitting the files. Rather, it appears that the transmission line disclosed in Srinivasan is a conventional dedicated line for providing mirrored storage.

Applicant respectfully submits that the above features are clearly absent from Srinivasan, and notes that the Office Action either misconstrues or takes out of context what is stated in Srinivasan to allege that the reference discloses such features. For example, the Examiner cites column 5, lines 17-36 as disclosing the interception of I/O transactions. However, this passage and the corresponding figure merely appear to exemplify that write transactions can be concurrently sent to the remote location over a dedicated line.

There is also no mention in this passage, or the other passages cited by the Examiner, of storing an exact copy of the write transactions in file system files, or transmitting such files using a non-proprietary network protocol to accommodate exact reproduction of the write

transactions at the remote location. For example, the Examiner apparently takes the position that Srinivasan makes an exact copy of the write transactions because Srinivasan “maintains a copy of the dataset.” (Office Action, p. 3). In Srinivasan, the dataset is the data itself (*e.g.*, database, file, etc.) that appears at the primary and secondary locations. This dataset may be revised, and the revisions to the dataset are managed to ensure consistency at both locations. However, the fact that the dataset is revised does not entail retention of the exact copy of write transactions as claimed.

Moreover, the related features of “*storing said exact copy within a series of files that are created on a file-system of the local storage system*” are clearly absent from Srinivasan. Here, the Examiner appears to state that “*i.e. dataset*” and storage of the dataset entail these features. (Office Action, p. 3). The dataset is not the write transactions, nor therefore an exact copy of the write transactions. The write transactions are separate entities, even if they are used to effect what the dataset becomes. Srinivasan thus does not disclose storing the exact copy of the write transactions as claimed.

Finally, in general, it is noted that the Examiner’s position is illogical, or at least logically inconsistent. This is because the Examiner has construed the dataset of Srinivasan as being not only the database or other data that may be stored on the disk, but also a copy of the write transactions that had been applied to that dataset, as well as an example of storing the exact copy of the write transactions in a series of files, and finally transmitting the same (*i.e.*, the dataset) to the secondary location. The “dataset” of Srinivasan cannot be reasonably construed as disclosing all of these features, and Applicant respectfully requests a coherent explanation as to how this possibly could be the case.

For reasons similar to those provided regarding claim 1, independent claims 3 and 5 are also neither disclosed nor suggested by Srinivasan. Additionally, dependent claims 1, 2, 4, 6, 9-11, 18-21, 23 and 25 respectively incorporate the features recited in the independent claims as well as their separately recited, patentably distinct features, and thus are not disclosed by Srinivasan as well.

Accordingly, Applicant respectfully requests reconsideration and withdrawal of the rejection of the noted claims under 35 U.S.C. § 102 as being anticipated by Srinivasan.

Claims 7 has been rejected under 35 U.S.C. § 103(a) as being unpatentable over Srinivasan in view U.S. Patent No. 5,673,382 to Cannon et al. ("Cannon"); and claim 8 has been rejected as being unpatentable over Srinivasan and further in view of U.S. Patent No. 5,713,014 to Durflinger et al. ("Durflinger"). These rejections are respectfully traversed.

Claims 7 and 8 incorporate the features of claim 1 described as being absent from Srinivasan. As stated in the previous response, Durflinger and Cannon clearly do not disclose the above-described features that are noted as being absent from Srinivasan. Particularly, the references fail to disclose or suggest retention of the exact copy of the series of write transactions, doing such in a series of file system based files, non-proprietary transmission of the series of files to the remote system, and exact reproduction of the series of write transactions on the remote storage system, all as claimed by Applicant.

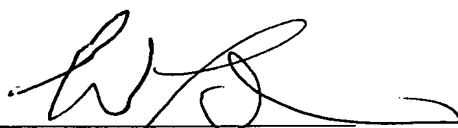
Moreover, claims 7 and 8 recite additional features for the claimed file system files. Since even the basic elements of such files are not found in the cited references, clearly there is also a failure to disclose the particulars recited in claims 7 and 8. Cannon appears to describe locating a file by noting its offset within a storage volume. (Cannon, 8:41-46). This does not disclose the presentation of a write transaction in a file, and corresponding inclusion of the size of the file itself, or offset information as claimed. Durflinger discloses a database management system that uses pointers to locate data positions within files. Usage of pointers for database management is clearly in a different context, and markedly different from Applicant's claimed invention.

Finally, Applicant notes that a proper motivation to combine the various relied upon references has not been made of record, which also supports a conclusion that a prima facie case of obviousness has not been presented by the Examiner.

Accordingly, Applicant requests reconsideration and withdrawal of the rejection of claims 7-8 under 35 U.S.C. § 103.

For the foregoing reasons, reconsideration and allowance of the claims which remain in the application are solicited. If any further issues remain, the Examiner is invited to telephone the undersigned to resolve them.

Date: February 15, 2007

By   
Christopher M. Tobin  
Registration No.: 40,290  
RADER, FISHMAN & GRAUER PLLC  
1233 20th Street, N.W., Suite 501  
Washington, DC 20036  
(202) 955-8779  
Attorney for Applicant

**RADER, FISHMAN & GRAUER, PLLC**  
Lion Building, 1233 20<sup>th</sup> Street, N.W., Suite 501  
Washington, D.C. 20036  
Tel: (202) 955-3750; Fax: (202) 955-3751